UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/014,742	10/30/2001	Beat Heeb	20296-300201	5753

7590	06/10/2005	

OPPENHEIMER WOLFF & DONNELLY, LLP
Suite 600
1620 L Street, NW
Washington, DC 20036

EXAMINER
NGUYEN BA, HOANG VU A

ART UNIT	PAPER NUMBER
2192	

DATE MAILED: 06/10/2005

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(application filed on or after May 29, 2000)

The Patent Term Adjustment to date is 675 day(s). If the issue fee is paid on the date that is three months after the mailing date of this notice and the patent issues on the Tuesday before the date that is 28 weeks (six and a half months) after the mailing date of this notice, the Patent Term Adjustment will be 675 day(s).

If a Continued Prosecution Application (CPA) was filed in the above-identified application, the filing date that determines Patent Term Adjustment is the filing date of the most recent CPA.

Applicant will be able to obtain more detailed information by accessing the Patent Application Information Retrieval (PAIR) WEB site (http://pair.uspto.gov).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571) 272-7702. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at (703) 305-8283.

RECEIVED
OIPE/IAP

JUN 2 1 2005

PTOL-85 (Rev. 12/04) Approved for use through 04/30/2007.

UNITED STATES PATENT AND TRADEMARK OFFICE

NOTICE OF ALLOWANCE AND FEE(S) DUE

7590	06/10/2005	

OPPENHEIMER WOLFF & DONNELLY, LLP
Suite 600
1620 L Street, NW
Washington, DC 20036

EXAMINER
NGUYEN BA, HOANG VU A

ART UNIT	PAPER NUMBER
2192	

DATE MAILED: 06/10/2005

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/014,742	10/30/2001	Beat Heeb	20296-300201	5753

TITLE OF INVENTION: METHOD FOR FAST COMPILATION OF PREVERIFIED JAVA TM BYTECODE TO HIGH QUALITY NATIVE MACHINE CODE

APPLN. TYPE	SMALL ENTITY	ISSUE FEE	PUBLICATION FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	NO	$1400	$300	$1700	09/12/2005

THE APPLICATION IDENTIFIED ABOVE HAS BEEN EXAMINED AND IS ALLOWED FOR ISSUANCE AS A PATENT. **PROSECUTION ON THE MERITS IS CLOSED.** THIS NOTICE OF ALLOWANCE IS NOT A GRANT OF PATENT RIGHTS. THIS APPLICATION IS SUBJECT TO WITHDRAWAL FROM ISSUE AT THE INITIATIVE OF THE OFFICE OR UPON PETITION BY THE APPLICANT. SEE 37 CFR 1.313 AND MPEP 1308.

THE ISSUE FEE AND PUBLICATION FEE (IF REQUIRED) MUST BE PAID WITHIN **THREE MONTHS** FROM THE MAILING DATE OF THIS NOTICE OR THIS APPLICATION SHALL BE REGARDED AS ABANDONED. **THIS STATUTORY PERIOD CANNOT BE EXTENDED.** SEE 35 U.S.C. 151. THE ISSUE FEE DUE INDICATED ABOVE REFLECTS A CREDIT FOR ANY PREVIOUSLY PAID ISSUE FEE APPLIED IN THIS APPLICATION. THE PTOL-85B (OR AN EQUIVALENT) MUST BE RETURNED WITHIN THIS PERIOD EVEN IF NO FEE IS DUE OR THE APPLICATION WILL BE REGARDED AS ABANDONED.

HOW TO REPLY TO THIS NOTICE:

I. Review the SMALL ENTITY status shown above.

If the SMALL ENTITY is shown as YES, verify your current SMALL ENTITY status:

A. If the status is the same, pay the TOTAL FEE(S) DUE shown above.

B. If the status above is to be removed, check box 5b on Part B - Fee(s) Transmittal and pay the PUBLICATION FEE (if required) and twice the amount of the ISSUE FEE shown above, or

If the SMALL ENTITY is shown as NO:

A. Pay TOTAL FEE(S) DUE shown above, or

B. If applicant claimed SMALL ENTITY status before, or is now claiming SMALL ENTITY status, check box 5a on Part B - Fee(s) Transmittal and pay the PUBLICATION FEE (if required) and 1/2 the ISSUE FEE shown above.

II. PART B - FEE(S) TRANSMITTAL should be completed and returned to the United States Patent and Trademark Office (USPTO) with your ISSUE FEE and PUBLICATION FEE (if required). Even if the fee(s) have already been paid, Part B - Fee(s) Transmittal should be completed and returned. If you are charging the fee(s) to your deposit account, section "4b" of Part B - Fee(s) Transmittal should be completed and an extra copy of the form should be submitted.

III. All communications regarding this application must give the application number. Please direct all communications prior to issuance to Mail Stop ISSUE FEE unless advised to the contrary.

IMPORTANT REMINDER: Utility patents issuing on applications filed on or after Dec. 12, 1980 may require payment of maintenance fees. It is patentee's responsibility to ensure timely payment of maintenance fees when due.

PTOL-85 (Rev. 12/04) Approved for use through 04/30/2007.

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), to: <u>Mail</u>

Mail Stop ISSUE FEE
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

or <u>Fax</u> (703) 746-4000

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)	
7590 06/10/2005 OPPENHEIMER WOLFF & DONNELLY, LLP Suite 600 1620 L Street, NW Washington, DC 20036	Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission. **Certificate of Mailing or Transmission** I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being facsimile transmitted to the USPTO (703) 746-4000, on the date indicated below.

_____ (Depositor's name)

_____ (Signature)

_____ (Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/014,742	10/30/2001	Beat Heeb	20296-300201	5753

TITLE OF INVENTION: METHOD FOR FAST COMPILATION OF PREVERIFIED JAVA TM BYTECODE TO HIGH QUALITY NATIVE MACHINE CODE

APPLN. TYPE	SMALL ENTITY	ISSUE FEE	PUBLICATION FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	NO	$1400	$300	$1700	09/12/2005

EXAMINER	ART UNIT	CLASS-SUBCLASS
NGUYEN BA, HOANG VU A	2192	717-154000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).

❏ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.

❏ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-02 or more recent) attached. **Use of a Customer Number is required.**

2. For printing on the patent front page, list

(1) the names of up to 3 registered patent attorneys or agents OR, alternatively,

(2) the name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.

1 _____

2 _____

3 _____

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document has been filed for recordation as set forth in 37 CFR 3.11. Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE   (B) RESIDENCE: (CITY and STATE OR COUNTRY)

Please check the appropriate assignee category or categories (will not be printed on the patent) :   ❏ Individual ❏ Corporation or other private group entity ❏ Government

4a. The following fee(s) are enclosed:

❏ Issue Fee

❏ Publication Fee (No small entity discount permitted)

❏ Advance Order - # of Copies _____

4b. Payment of Fee(s):

❏ A check in the amount of the fee(s) is enclosed.

❏ Payment by credit card. Form PTO-2038 is attached.

❏ The Director is hereby authorized by charge the required fee(s), or credit any overpayment, to Deposit Account Number _____ (enclose an extra copy of this form).

5. Change in Entity Status (from status indicated above)

❏ a. Applicant claims SMALL ENTITY status. See 37 CFR 1.27.   ❏ b. Applicant is no longer claiming SMALL ENTITY status. See 37 CFR 1.27(g)(2).

The Director of the USPTO is requested to apply the Issue Fee and Publication Fee (if any) or to re-apply any previously paid issue fee to the application identified above.
NOTE: The Issue Fee and Publication Fee (if required) will not be accepted from anyone other than the applicant; a registered attorney or agent; or the assignee or other party in interest as shown by the records of the United States Patent and Trademark Office.

Authorized Signature _____   Date _____

Typed or printed name _____   Registration No. _____

PTOL-85 (Rev. 12/04) Approved for use through 04/30/2007.   OMB 0651-0033   U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

	Application No.	**Applicant(s)**
Notice of Allowability	10/014,742	HEEB, BEAT
	Examiner	**Art Unit**	
	Hoang-Vu A. Nguyen-Ba	2192	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☒ This communication is responsive to _Amendment filed December 13, 2004_.

2. ☒ The allowed claim(s) is/are _1-18, 20-22 and 24_.

3. ☐ The drawings filed on _____ are accepted by the Examiner.

4. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

    a) ☐ All   b) ☐ Some*   c) ☐ None of the:

    1. ☐ Certified copies of the priority documents have been received.

    2. ☐ Certified copies of the priority documents have been received in Application No. _____ .

    3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

    * Certified copies not received: _____ .

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. ☐ A SUBSTITUTE OATH OR DECLARATION must be submitted. Note the attached EXAMINER'S AMENDMENT or NOTICE OF INFORMAL PATENT APPLICATION (PTO-152) which gives reason(s) why the oath or declaration is deficient.

6. ☒ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.

    (a) ☐ including changes required by the Notice of Draftsperson's Patent Drawing Review (PTO-948) attached

        1) ☐ hereto or 2) ☐ to Paper No./Mail Date _____ .

    (b) ☒ including changes required by the attached Examiner's Amendment / Comment or in the Office action of

    Paper No./Mail Date _____ :

    Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).

7. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

1. ☒ Notice of References Cited (PTO-892)

2. ☐ Notice of Draftperson's Patent Drawing Review (PTO-948)

3. ☒ Information Disclosure Statements (PTO-1449 or PTO/SB/08), Paper No./Mail Date _11/22/04_

4. ☐ Examiner's Comment Regarding Requirement for Deposit of Biological Material

5. ☐ Notice of Informal Patent Application (PTO-152)

6. ☐ Interview Summary (PTO-413), Paper No./Mail Date _____ .

7. ☒ Examiner's Amendment/Comment

8. ☒ Examiner's Statement of Reasons for Allowance

9. ☒ Other _App'd Drawings&Fax Sheet_.

ANTONY NGUYEN-BA
PRIMARY EXAMINER

Continuation of Attachment(s) 9. Other: Approved Drawings & Fax Cover Sheet.

EXAMINER'S AMENDMENT and
EXAMINER'S STATEMENT OF REASON FOR ALLOWANCE

1.　　This action is responsive to amendment filed December 13, 2004.

Response to Amendments

2.　　Per applicant's request, claims 1-19 have been amended; new claims 20-24 have been added. Claims 1-24 are pending.

3.　　The objection to the drawings made in the previous office action is withdrawn in view of Applicants' amendments to the drawings to correct the identified informalities. Copies of the approved corrected drawings with the Examiner's initial are attached to this Office action. However, corrected drawings **labeled** as "Replacement Sheet(s)" must be submitted.

INFORMATION ON HOW TO EFFECT DRAWING CHANGES

Replacement Drawing Sheets

Drawing changes must be made by presenting replacement sheets which incorporate the desired changes and which comply with 37 CFR 1.84. An explanation of the changes made must be presented either in the drawing amendments section, or remarks, section of the amendment paper. Each drawing sheet submitted after the filing date of an application must be labeled in the top margin as either "Replacement Sheet" or "New Sheet" pursuant to 37 CFR 1.121(d). A replacement sheet must include all of the figures appearing on the immediate prior version of the sheet, even if only one figure is being amended. The figure or figure number of the amended drawing(s) must not be labeled as "amended." If the changes to the drawing figure(s) are not accepted by the examiner, applicant will be notified of any required corrective

action in the next Office action. No further drawing submission will be required, unless applicant is notified.

Identifying indicia, if provided, should include the title of the invention, inventor's name, and application number, or docket number (if any) if an application number has not been assigned to the application. If this information is provided, it must be placed on the front of each sheet and within the top margin.

Timing of Corrections

Applicant is required to submit acceptable corrected drawings within the time period set in the Office action. See 37 CFR 1.85(a). Failure to take corrective action within the set period will result in ABANDONMENT of the application.

If corrected drawings are required in a Notice of Allowability (PTOL-37), the new drawings MUST be filed within the THREE MONTH shortened statutory period set for reply in the "Notice of Allowability." Extensions of time may NOT be obtained under the provisions of 37 CFR 1.136 for filing the corrected drawings after the mailing of a Notice of Allowability.

4.   The objection to the specification is withdrawn in view of Applicants' amendments to the specification to correct minor informalities.

5.   The objection to claims 2, 6 and 9 is withdrawn in view of Applicants' amendments to these claims to correct minor informalities.

6.   The rejection of claims 3-19 under 35 U.S.C. § 112, first paragraph, as failing to comply with the written description requirement is withdrawn in view of Applicants' amendments to the specification to clarify the limitations "using information from preceding instructions to mimic an optimizing compiler" and "to mimic an optimizing compiler" recited in the claims.

7.   The rejection of claims 2-19 under 35 U.S.C. § 112, second paragraph, as being indefinite is withdrawn in view of Applicants' amendments to these claims to provide proper antecedent basis to the identified terms.

8.   The provisional obviousness-type double patenting rejection of claims 1-19 over claims 1-29 of copending Application No. 10/016,794 is withdrawn in view of Applicants' filing of a terminal disclaimer.

9.   The rejection of claims 1 and 3-8 under 35 U.S.C. § 102(a) is withdrawn in view of Applicants' amendments to these claims to incorporate the identified allowable subject matter.

Examiner's Amendment

10.   An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

Authorization for this examiner's amendment was given in a telephone interview with William F Ahmann, Reg. No. 52,548 on May 20, May 26, and June 3, 2005.

The application has been amended as follows:

a.   **Claim 3:**

i.   in line 11, after "said native machine code", delete the period and insert -- ; and --

ii.   in line 12, insert -- producing said native machine code in a single sequential pass in which information from preceding instruction translation is used to perform the same optimizing process of an optimizing compiler without the extensive memory and time requirements. --

b.   Cancel **Claim 19**

c.   **Claim 9:** in line 5, after "setting" delete "said"

d.   **Claim 13:** in line 4, after "method duplicating or reordering" delete "said" and insert — a —

e.   **Claim 14:** in line 4, after "emitting native machine code using", delete "said"

f.   **Claim 15:** in line 4, before "stack mapping information", delete — said --

g.   **Claim 21:** replace claim 21 (page 15 of Amendments to the Claims) with the following:

21.   (currently amended)   A computer implemented method, comprising:

processing a first bytecode of a sequence of bytecodes;

processing a second bytecode of the sequence of bytecodes using information associated with resulting from the processing of the first bytecode; and

producing optimized native machine code in a single pass through the sequence of the bytecodes, using preceding translation information to optimize the native machine code.

h.   Cancel **Claim 23**.

Examiner's Statement of Reasons for Allowance

11.   Claims 1-18, 20-22 and 24 are allowed.

12.   The following is an examiner's statement of reasons for allowance.

The prior art of record, i.e., applicant's admitted prior art (APA), teaches attempts to improve compilation speed by compilation during idle times and by pre-verification. However, APA fails to teach or suggest improving compilation speed using a second code segment that creates optimized machine code from bytecodes received by a first code segment in a **single** sequential **pass** in which **information**

from preceding instruction translation is used to perform the same optimizing process of an optimized compiler without the extensive memory and time requirement, as recited in the amended independent claims 1, 3 and 20.

Furthermore, APA does not teach or suggest the limitations recited in claim 2.

Moreover, a copy of a newly found reference, which is now made of record, i.e., Azevedo-Nicolau-Hummel, Java™ Annotation-Aware Just-in-time (AJIT) Compilation System ("AJIT") was faxed to Applicants on June 1, 2005 for review. On June 2 and 3, 2005, Applicants' representative and the examiner discussed over the telephone the claim language of Claim 21 in light of the teachings of AJIT. Agreement was reached that by amending the claim language to read "using information resulting from the processing of the first bytecode" instead of "using information associated with the processing of the first bytecode," Claim 21 would distinguish over the teachings of AJIT.

AJIT presents an alternative to an optimizing JIT compiler that makes use of code annotations generated by a Java™ front-end (i.e., Java™ compiler, a.k.a., Javac). These annotations carry information concerning compiler optimization. During the translation process, an annotation-aware JIT (AJIT) system then uses this information to produce high-performance native code without performing much of the necessary analyses or transformations. The process of producing native code from an annotated Java™ bytecodes is done in a single pass over the bytecode stream. As each bytecode and its annotations bytes are read, a corresponding Kaffe IR operation is generated. The generated Kaffe IR operation depends on the information provided by the annotations. This information may suggest that the bytecode translation be entirely skipped, or that some sub-operations be eliminated or simplified. AJIT, however, fails to teach or suggest that processing a second bytecode of the sequence of bytecodes using **information resulting from the processing of the first**

bytecode and producing optimized native machine code in a single pass through the sequence of bytecodes, **using preceding translation information** to optimize the native machine code (instant claims 21, 1 and 3). The Kaffe IR operation in AJIT depends on the information provided by the annotations that are **associated with** the bytecode that is **currently** processed as opposed to the information **resulting from** the processing of the **preceding** bytecode (emphasis added) as claimed in the present invention.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submission should be clearly labeled "Comments on Statement of Reasons for Allowance."

13.   Any inquiry concerning this communication or earlier communications from the examiner should be directed to Hoang-Vu "Antony" Nguyen-Ba whose telephone number is (571) 272-3701. The Examiner can normally be reached on Tuesday-Friday, 7:15 to 17:15.

If attempts to reach the Examiner by telephone are unsuccessful, the Examiner's supervisor, Tuan Dam can be reached at (571) 272-3695. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov.

Should you have questions on access to the Private PAIR system, contact the

Electronic Business Center (EBC) at 866-217-9197 (toll-free).

ANTONY NGUYEN-BA
PRIMARY EXAMINER

Art Unit 2192

June 8, 2005

				COMPLETE IF KNOWN	
INFORMATION DISCLOSURE STATEMENT BY APPLICANT Form PTO-1449 (Modified) (Use several sheets if necessary)			Application Number	10/014,742	
			Confirmation Number	5753	
			Filing Date	October 30, 2001	
			First Named Inventor	Beat Heab	
			Group Art Unit	~~2123~~ 2192	
			Examiner Name	H.A. Nguyen-Ba	
Sheet	1	of	1	Attorney Docket No.	59296-8002.US01

U.S. PATENT DOCUMENTS

Examiner Initials	Cite No.	NUMBER	Kind Code (if known)	Name of Patentee or Inventor of Cited Document	Date of Publication or Filing Date of Cited Document	Pages, Columns, Lines, Where Relevant Figures Appear

FOREIGN PATENT DOCUMENTS

Examiner Initial	Cite No.	Office	NUMBER	Kind Code (if known)	Name of Patentee or Applicant of Cited Document	Date of Publication or Filing Date of Cited Document	Pages, Columns, Lines, Where Relevant Figures Appear	T

OTHER PRIOR ART-NON PATENT LITERATURE DOCUMENTS

Examiner Initials	Cite No.	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc.), date, page(s), volume issue number(s), publisher, city and/or country where published.	T
(initialed)		Alpern et al., "The Jalapeño Virtual Machine," IBM Systems Journal, Vol. 39, No. 1, 2000, pp. 211-238.	

EXAMINER *(signature)*	DATE CONSIDERED 5/19/05

*EXAMINER: Initial if reference considered, whether or not criteria is in conformance with MPEP 609. Draw line through citation if not in conformance and not considered. Include copy of this form with next communication to application(s).

BY043210.130

UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

Fax Cover Sheet

Date: 01 Jun 2005

To: William F. Ahmann	**From:** Hoang-Vu A. Nguyen-Ba
Application/Control Number: 10/014,742	**Art Unit:** 2192
Fax No.: (650) 838-4350	**Phone No.:** (571) 272-3701
Voice No.: (650) 838-4300	**Return Fax No.:** (703) 872-9306
Re: Claims 1, 3, 21 in view of attached reference	**CC:**

☐ **Urgent**   ☑ **For Review**   ☐ **For Comment**   ☐ **For Reply**   ☐ **Per Your Request**

Comments:
Please see section 3, Annotation-Aware JIT Compilation System, 3rd paragraph of the attached reference

Number of pages 11 including this page

Title: METHOD FOR FAST COMPILATION OF
PREVERIFIED JAVA TM BYTE CODE TO HIGH
QUALITY NATIVE MACHINE CODE
Inventor: Beat Heeb   REPLACMENT SHEET
Docket No.: 59296-8002.US01
Perkins Cole LLP (650) 838-4300

100 ⟶   1/5

Approved by examiner Hen 5/28/05

102

Loop through byte-code instructions;

Analyze one byte-code instruction;

Is the byte-code instruction the last?

Y ⟶ 108

End

N

104

Determine stack status from byte-code
instruction and store in stack status
storage

FIG. 1A
(Prior Art)

Title: METHOD FOR FAST COMPILATION OF
PREVERIFIED JAVA TM BYTE CODE TO HIGH
QUALITY NATIVE MACHINE CODE
Inventor: Beat Heeb   REPLACMENT SHEET
Docket No.: 59296-8002.US01
Perkins Coie LLP (650) 838-4300

150 ➔   2/5

Approved by Examiner 9/cm 5/28/05

Loop through byte-code instructions;

Analyze one byte-code instruction;

Is the byte-code instruction the last?

152

Y   158

End

N

154

Use stack status storage and byte-code instruction to translate to machine code

FIG. 1B
(Prior Art)

Title: METHOD FOR FAST COMPILATION OF
PREVERIFIED JAVA TM BYTE CODE TO HIGH
QUALITY NATIVE MACHINE CODE
Inventor: Beat Heeb   REPLACMENT SHEET
Docket No.: 59296-6002.US01
Perkins Coie LLP (650) 838-4300

200 ———►   3/5

Approved by examiner Yau 3/28/05

```
                                    ┌─────────────┐
                                    │    Start    │
                                    └──────┬──────┘
                                           │                    ╭─ 202
                                           ▼
                            ┌──────────────────────────────┐
     ┌─ 248                 │   Select first class to compile │
┌───┐   ┌──────────────┐    └──────────────┬─────────────────┘
│ D ├──►│ Select next  ├───────────────────┤              ╭─ 204
└───┘   │ class to     │                   ▼
        │ compile      │    ┌──────────────────────────────┐
        └──────────────┘    │      Select first method       │
                            └──────────────┬─────────────────┘
     ┌───┐ ┌──────────────┐                │              ╭─ 205
     │ C ├►│ Select next  ├────────────────┤
     └───┘ │ method to    │                ▼
           │ compile      │  ┌──────────────────────────────┐
           └──────┬───────┘  │ Create storage to store actual │
                  └─ 246      │ mappings and native code       │
                             │ addresses                      │
                             └──────────────┬─────────────────┘
                                            │              ╭─ 206
                                            ▼
                             ┌──────────────────────────────┐
                             │ Initialize stack mappings to   │
                             │ "empty" and addresses to       │
                             │ "unknown"                      │
                             └──────────────┬─────────────────┘
                                            │              ╭─ 207
                                            ▼
                             ┌──────────────────────────────┐
     ┌─ 244                  │    Select first instruction    │
┌───┐   ┌──────────────┐     └──────────────┬─────────────────┘
│ B ├──►│ Select next  ├────────────────────┤
└───┘   │ instruction  │                    ▼          ╭─ 208
        │ to compile   │            ╱◆◆◆◆◆◆◆◆◆◆◆◆◆◆╲
        └──────────────┘     Y   ╱   Is a stack      ╲
                          ◄──────  mapping stored for ─
                                  ╲ the actual         ╱
                                   ╲ instruction?     ╱
                                    ╲◆◆◆◆◆◆◆◆◆◆◆◆◆╱
                                           │ N
                                           ▼          ╭─ 212
```

Emit code to store stacks and set stack mapping to "stack"   ── 212

Is direct control flow present?   ── 210

Read stack layout from stack mapping and set mapping to "stack"   ── 214

Set native code address for actual instruction   ── 216

```
                                    ┌───┐
                                    │ A │
                                    └───┘
```

FIG. 2A

Title: METHOD FOR FAST COMPILATION OF
PREVERIFIED JAVA TM BYTE CODE TO HIGH
QUALITY NATIVE MACHINE CODE
Inventor: Beat Heeb   REPLACMENT SHEET
Docket No.: 59298-8002.US01
Perkins Coie LLP (650) 838-4300

200 → (A)   4/5

Approved by Examiner 1601 5/28/05

218 Is actual instruction "load constant"? — Y → **220** Create new constant stack mapping

N

222 Is actual instruction "load local"? — Y → **224** Create new local stack mapping

N

226 Is actual instruction a stack manipulating instruction? — Y → **228** Duplicate and/or reorder stack mapping according to instruction

N

230 Is actual instruction a jump or switch instruction? — Y → **232** Emit code using stack mapping information and emit code to store unused stack values

N

234 Is actual instruction any other instruction? — Y → **236** Emit code using stack mapping information

N

242 Next class? — N → End
242 Next class? — Y → D

240 Next method? — N → 242
240 Next method? — Y → C

238 Next instruction? — N → 240
238 Next instruction? — Y → B

FIG. 2B

Title: METHOD FOR FAST COMPILATION OF
PREVERIFIED JAVA TM BYTE CODE TO HIGH
QUALITY NATIVE MACHINE CODE
Inventor: Beat Heeb   REPLACMENT SHEET
Docket No.: 59298-8002.US01
Perkins Coie LLP (650) 838-4300

300 ⟶

5/5

Approved by examiner Han 5/28/05

For each value on the bytecode stack	A field showing actual mapping to storage in target machine	Constant
		Local
		Temporary
		Stack
	A field containing additional information	Constant Value
		Slot Number
		Register Number
For each target of a jump or switch instruction	A field to store native code address	

FIG. 3

		Application/Control No.	Applicant(s)/Patent Under Reexamination
Notice of References Cited		10/014,742	HEEB, BEAT
		Examiner	Art Unit	Page 1 of 1
		Hoang-Vu A. Nguyen-Ba	2192	

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
	A	US-			
	B	US-			
	C	US-			
	D	US-			
	E	US-			
	F	US-			
	G	US-			
	H	US-			
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	Azevedo-Nicolau-Hummel, Java Annotation-Aware Just-In-Time (AJIT) Compilation System, 06/1999, http://www.cs.ucsb.edu/conferences/java99/papers/63-azevedo.pdf
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Java Annotation-Aware Just-In-Time (AJIT) Compilation System

Ana Azevedo,* Alex Nicolau

University of California, Irvine
aazevedo, nicolau@ics.uci.edu

Joe Hummel

University of Illinois, Chicago
jhummel@eecs.uic.edu

Abstract

The Java Bytecode language lacks expressiveness for
traditional compiler optimizations, making this portable,
secure software distribution format inefficient as a pro-
gram representation for high performance. This ineffi-
ciency results from the underlying stack model, as well
as the fact that many bytecode operations intrinsically
include sub-operations (e.g., `iload` includes the ad-
dress computation, array bounds checks and the actual
load of the array element). The stack model, with no
operand registers and limiting access to the top of the
stack, prevents the reuse of values and bytecode reorder-
ing. In addition, the bytecodes have no mechanism to
indicate which sub-operations in the bytecode stream
are redundant or subsumed by previous ones. As a con-
sequence, the Java Bytecode language inhibits the ex-
pression of important compiler optimizations, including
common sub-expression elimination, register allocation
and instruction scheduling.

The bytecode stream generated by the Java front-
end is a significantly under-optimized program repre-
sentation. The most common solution to overcome this
aspect of the language is the use of a Just-in-Time (JIT)
compiler to not only generate native code, but perform
optimization as well. However, the latter is a time con-
suming operation in an already time-constrained trans-
lation process. In this paper we present an alternative
to an optimizing JIT compiler that makes use of code
annotations generated by the Java front-end. These an-
notations carry information concerning compiler opti-
mization. During the translation process, an annotation-
aware JIT (AJIT) system then uses this information
to produce high-performance native code without per-

forming much of the necessary analyses or transforma-
tions. We describe the implementation of the first pro-
totype of our annotation-aware JIT system and show
performance results comparing our system with other
Java Virtual Machines (JVMs) running on SPARC ar-
chitecture.

1 Introduction

The Java Bytecodes are emerging as a software distri-
bution language for both its portability and safety fea-
tures. The portability property of the language is en-
sured by the platform-independent stack machine model
targeted by Java compilers. On the target machine,
this intermediate code representation is either inter-
preted [17], or compiled into native code using tradi-
tional ahead-of-time [14] or just-in-time compilers [1,
2, 16, 18, 24]. The safety features of the language are
based on the security violation checks performed at load
and run-time [11]. Such checks include enforcement
of methods and variables access modifiers, strict type-
checking and array bounds checking. Many of these
checks are implicit in the bytecodes, forcing the JVM
to perform them unless it can prove at load-time (via
analysis) that the checks are unnecessary.

In the design of the Java Bytecode language, a great
deal of effort was spent to make it secure and portable.
However, in order to be widely accepted, it must also
yield efficient execution on a wide range of machine ar-
chitectures. Unfortunately this is the weakest aspect
of the language and is currently the focus of much re-
search. The inefficient execution of Java Bytecode pro-
grams lies with the definition of the bytecodes them-
selves. The language is poor for conveying the result of
many common and important compiler optimizations
that are traditionally expressed in the native code gen-
erated by optimizing compilers. The direct translation
of a bytecode stream generated by a Java front-end into
target machine code results in low-quality code.

The first limitation in expressing compiler optimiza-
tion is the stack model of the Java Bytecodes. This

Java Code
| ```
public static void foo(int a[], int b[], int offset1, int offset2){
    for (int i=0; i<a.length; i++)
        a[i] = b[i] + offset1 + offset2;
}
``` |

IR	Optimized IR	Optimized Bytecode
| ```
1 : smovi 0, i
2 : aadd a, "arraySizeOffset", T1
3 : ild (T1), T2
4 : icmpge i, T2, T3
5 : br T3 (18)
6 : ishl i, "ishift", T5
7 : iadd T5, "arraySizeOffset", T6
8 : aadd b, T6, T7
9 : ild (T7), T4
10 : iadd T4, offset1, T8
11 : iadd T8, offset2, T9
12 : ishl i, "ishift", T10
13 : iadd T10, "arraySizeOffset", T11
14 : aadd a, T11, T12
15 : ist  T9, (T12)
16 : iadd  i, 1 , i
17 : jmp (2)
18 : return
``` | ```
1 : iadd offset1, offset2, T1
2 : smovi 0 , i
3 : aadd a, "arraySizeOffset", T2
4 : ild (T2), T3
5 : icmpge i, T3, T4
6 : br T4 (16)
7 : ishl i, "ishift", T6
8 : iadd T6, "arraySizeOffset", T7
9 : aadd b, T7, T8
10 : ild (T8), T5
11 : iadd T5, T1, T9
12 : aadd a, T7, T10
13 : ist T9, (T10)
14 : iadd i, 1 , i
15 : jmp (5)
16 : return
``` | ```
0  iload_2
1  iload_3
2  iadd
3  istore 5
5  aload_0
6  arraylength
7  istore 6
9  iconst_0
10 istore 4
12 goto 29
15 aload_0
16 iload 4
18 aload_1
19 iload 4
21 iaload
22 iload 5
24 iadd
25 iastore
26 iinc 4 1
29 iload 4
31 iload 6
33 if_icmplt 15
36 return
``` |

Figure 1: Java Bytecodes as a language for program representation

model provides no registers and restricts access to only the top element of the stack. Restricting access to the top of the stack prevents the reordering of bytecodes, a necessary transformation during instruction scheduling. And without registers to hold values, the stack model sequentializes computation and prevents the reuse of values (since again, only the top is accessible). Obviously, the lack of registers also prevents the expression of register allocation, a critical and potentially time-consuming optimization.

The second limitation of the Java Bytecodes as a program representation is the fact that many bytecodes intrinsically encapsulate many machine sub-operations (e.g., iaload includes the address computation, array bounds checks and the actual load of the array element). The Java front-end can detect when sub-operations are redundant or subsumed by preceding sub-operations, and can try to apply traditional code-improving transformations in order to eliminate these sub-operations. However, the compiler is still limited by the stack-based nature of the Java Bytecodes, in which sub-operations cannot easily be separated, eliminated or rearranged. Furthermore, there is no mechanism in the language to disable sub-operations when deemed unnecessary. For this reason, straighforward compiler optimizations such as common sub-expression elimination, array bounds check elimination and loop-invariant code removal have limited expressiveness in Java Bytecodes.

To demonstrate these limitations, consider the example in Figure 1. This example assumes that a RISC-like, three address code Intermediate Representation (IR) is used in the Java front-end to bytecode com-

piler. The leftmost column shows the unoptimized IR[1] corresponding to the Java code at the top of Figure 1. The middle column shows the result of performing some simple optimizations, such as loop invariant removal of expression offset1 + offset2 and the array size reference. After optimizing this IR, the compiler is then able to produce the optimized bytecode stream shown in the last column. However, additional optimizations are possible that cannot be expressed in the final bytecode. For example, the sub-operations comprising array element accesses represent common sub-expressions and thus one could be eliminated (the index is the same for accessing the integer arrays a and b and therefore the array index computation in lines 6-7 and 12-13 in the leftmost column are redundant). Likewise, given the bounds on the loop, all array bounds checks involving a are unnecessary (and those involving b could be reduced to a single check before the loop starts). Clearly, the resulting bytecode has room for improvement.

The implication is that even though the Java front-end can compile a program into a clean and optimized sequence of bytecodes, a JIT compiler will still need to perform significant optimization in order to generate high-quality native code. This in turn implies that a JIT compiler will have to perform bytecode analysis to extract information about the program for the purposes of optimization. This introduces a potentially significant overhead in an already time-constrained JIT system. In this paper we present an alternative to the traditional optimizing JIT compiler based on bytecode annotations. In our annotation-aware JIT (AJIT) compilation system, the translation of bytecodes into

---

[1] Array bound checks have been omitted.

high-performance native code is accomplished with the help of extra analysis information carried along with the bytecodes in the form of annotations. Our idea of Java Bytecode annotations was first introduced in [15]; in this paper we present the details of the implementation our annotation-aware JIT system. In particular, we show how annotations are effective in carrying information concerning register allocation, common sub-expressions and value propagation. We also present some initial results on the performance of the code generated by our AJIT system, demonstrating that our approach outperforms other JVM implementations on the SPARC architecture.

The format of this paper is as follows. In the next section we present the structure of our annotation generating Java front-end and discuss the types and formats of the annotations implemented in our first prototype. We also provide details concerning our compile-time register allocator that produces the annotations in support of dynamic register allocation. In Section 3 we discuss our annotation-aware JIT (AJIT) system and show how it uses annotations to implement run-time register allocation and produce native code. Next, in Section 4 we discuss related work. Finally, in Section 5 we present some preliminary results on the performance of our AJIT system, followed by our conclusions and discussion of future work in Section 6.

2 Annotation-Generating Compilation System

The idea of annotating a program representation with analysis information produced by a front-end compiler stems from the need to reduce the workload of run-time code optimizing systems. We have chosen to annotate the Java Bytecode representation, given its commerical success and widespread availability induced by its write-once-run-anywhere capability. However, the concept of annotations is a general one, and thus can be applied to any program representation.

Our annotation types and formats vary with the kind of information that needs to be conveyed to the run-time code optimizing system. For example, it may consist of high-level program information that is not expressible in the lower-level program representation, or compile-time analysis information that is too time consuming to produce at run-time. Figure 2 gives an overview of a general annotation-generating compilation system with a number of different annotations that we are currently working on. During the initial Java to bytecode translation, our annotation-generating compiler behaves as a traditional compiler. It builds a three-address code intermediate representation flexible enough to represent all the sub-operations that form each bytecode. On this IR traditional code-improving

techniques (e.g., copy propagation, common sub expression elimination, loop invariant code removal and register allocation) are applied and an optimized IR produced. Once this stage has been reached, each operation (or sequence of operations) is translated into an optimized stream of Java Bytecodes. Next, an annotation generator also reads the optimized IR, along with the data provided by various compiler analyses, and produces a set of annotations. Finally, the compiler performs a mapping phase in which the bytecode operations are paired with their corresponding IR operations and annotations, and then stores the annotated bytecode into the appropriate class file.

For example, in the case of the Virtual Register Allocation (VRA) annotations (to be explained shortly), each bytecode is annotated with the source and destination registers allocated to the operands of that Java IR operation. Then, the bytecode stream is copied into the code attribute section of the class file together with the annotations, the latter being stored as an extra code attribute. Storing annotations in this way guarantees backward compatibility with existing JVMs, which by definition must ignore unknown code attributes [11].
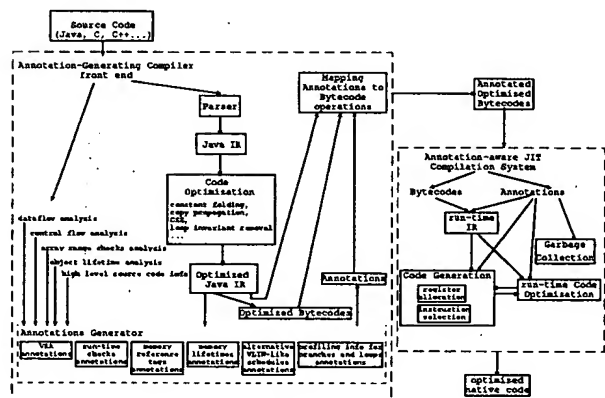

Figure 2: Annotation generating compiler and annotation aware JIT (AJIT) system

Our annotation-generating compiler was built on the freely available Java Bytecode compiler *guavac* version 0.3.1 [22]. From the Java source code, this compiler generates a parse tree and produces bytecodes. We augmented the compiler by (a) introducing functions for building and manipulating a three-address code IR, (b) implementing compiler optimizations for common sub-expression elimination, copy propagation and virtual register allocation, and (c) designing a Virtual Register Allocation annotation generator. This paper focuses in particular on the VRA annotations. The remaining annotations, as presented in Figure 2, are discussed in [15] or represent future work (see Section 6).

Virtual Register Allocation annotations represent the result of performing register allocation assuming an infinite number of *virtual* registers. The information provided by the VRA annotations is then used by the JIT compiler to perform a fast and efficient dynamic register allocation and also to indicate which bytecodes (or bytecode sub-operations) are redundant [2] or subsumed by preceding operations; such operations need not be translated into native code. How the JIT compiler interprets these annotations, does register allocation, and produces native code is explained in Section 3. In the remainder of this section we discuss the format of VRA annotations and how the Java front-end compiler produces them.

Each instruction defined in the Java Bytecode language is mapped into operations in our Java IR. Annotations for virtual register allocation basically hold information on the operands of the Java IR operations. The VRA annotations represent source operands, destination operands, and any intermediate values implicitly calculated by the bytecode sub-operations (e.g., array index calculation in an array load operation). For each bytecode instruction one or more VRA annotation formats exist. Each format indicates how a particular bytecode sub-operation should be translated: where to read its input operands, where to write the result, and perhaps whether or not this operation should be skipped entirely (e.g. when a previous operation has already computed the needed value).

Figure 3 shows an example of correspondence between bytecodes, Java IR and VRA annotations formats. Each SRC, EXTRA and DEST fields hold virtual register numbers representing the operands for the sub-operations. In Case 1 of Figure 3, the Java IR code sequence for the computation performed by the bytecode iaload is illustrated. The most general format of an iaload operation includes 2 SRC fields, 2 EXTRA fields and one DEST field with SRC-SRC-EXTRA-EXTRA-DEST as annotation header format. The first SRC field represents the virtual register that holds the array object reference; the second SRC field represents the virtual register that holds the index; the first EXTRA field represents the result of the array index calculation; the last EXTRA field represents the result of the array address calculation; and the DEST field represents the virtual register holding the array element read from memory. If the address computation has already been computed, as in Figure 3 Case 2, the header SRC-DEST indicates that the SRC field holds the array element address and DEST field is the suggested virtual register to hold the value read from memory, meaning that the translation process can skip the sub-operations for array index and address calculation and the bytecode iaload can be translated into a single load operation.

Case 1: Array element address calculation and array load

Bytecode	Java IR
iaload	V0 holds array address V1 holds index 1 : ishl V1, "ishift", V2 2 : iadd V2, "arraySizeOffset", V2 3 : aadd V0, V2, V3 4 : ild (V3), V4

Annotated Bytecode					
opcode	SRC	SRC	EXTRA	EXTRA	DEST
iaload	V0	V1	V2	V3	V4

Case 2: Array load

Bytecode	Java IR
iaload	V0 holds array element address 4 : ild (V0), V1

Annotated Bytecode		
opcode	SRC	DEST
iaload	V0	V1

Figure 3: Example of VRA annotations for `iaload` operation

In Figure 4, we show how local variables and class member variables are represented in our Java IR. Local variables are directly mapped to virtual registers. Local variable accesses (e.g, `iload` and `istore`) are represented in our Java IR as nop operations or move operations, annotated as SRC-DEST, CONST-DEST, CONST or SRC, depending on the result of optimizing the Java IR via copy propagation. When the JIT interprets the annotation formats SRC or CONST, it has the information that either (a) the local variable is in a virtual register indicated by the byte following the format header, or (b) it is a constant. In both cases, no machine code is generated for the bytecode. Class member variables are kept as variables in memory in our front-end compiler and accessed via load and store operations, as shown in Figure 4 for bytecodes `getstatic` and `getfield`. As a consequence, these variables are also kept in memory in our AJIT system. To enable some optimization on accesses to class member variables, we devised annotations that make explicit the variable address calculation, just like those in array references. For example, bytecode `getfield` has the different annotation formats SRC-DEST and EXTRA-EXTRA-EXTRA-DEST which state whether or not the variable's address has already been computed.

The choice of which virtual register to hold an operation's operands is crucial to the register allocation done at run-time. In order to enable a fast and efficient dynamic register allocation, the VRA annotations must convey the order in which variables should be allocated to physical registers (and thus which should be spilled if necessary). This is accomplished by assigning, at compile-time, the lowest virtual register numbers to the most important variables in the code. Then, at run-time, the register allocator should assign the lowest

---

single load operation.

---

[2] As discussed earlier, redundant bytecodes appear in the optimized bytecode stream due to the stack machine model.

Bytecode	Java IR	VRA Annotation Formats
iload	nop	SRC
istore	imov V1, V2	SRC DEST
	imov const, V1	CONST DEST
	nop	SRC
getstatic	amovi "addressOfClassField", V1 (b,c,s,i,l,d,f,a)ld (V1), V2	EXTRA DEST
	(b,c,s,i,l,d,f,a)ld (V1), V2	SRC DEST
	nop	SRC
getfield	amovi "offsetOfField", V2 aadd V1, V2, V3 (b,c,s,i,l,d,f,a)ld (V3), V4	SRC EXTRA EXTRA DEST
	amovi "offsetOfField", V2 aadd V1, V2, V3 (b,c,s,i,l,d,f,a)ld (V3), V4	SRC EXTRA EXTRA DEST
	aadd V1, V2, V3 (b,c,s,i,l,d,f,a)ld (V3), V4	SRC SRC EXTRA DEST
	(b,c,s,i,l,d,f,a)ld (V1), V2	SRC DEST
	nop	SRC

Figure 4: Example of VRA annotations for local variables and class member variables accesses

virtual register numbers to the physical machine registers. The details of our compile-time register allocation algorithm are presented in Section 2.1.

When designing the VRA annotations we opted for a format that was easy for the run-time system to decode so that processing the annotations would incur minimal overhead. The general VRA annotation formats include a byte-long header followed by a variable number of bytes representing the virtual register numbers. The header indicates how the subsequent annotation bytes should be interpreted. In our first prototype, we did not try to optimize the space consumed by the annotations, and thus we found that our annotations can double the size of the bytecode stream [15]. Another, potentially more significant problem is that of verification: to maintain security, a scheme is needed to verify the safety of the annotations in the class file, as malicious or incorrect annotations can lead to unsafe native code. We have not yet devised an algorithm for validating VRA annotations.

2.1 Compile-time Register Allocation

In our annotation-generating compiler we implement a modified priority-based graph-coloring algorithm. In a traditional Chaitin-style graph coloring algorithm [4, 5], an interference graph is pruned to decide the ordering in which live ranges are assigned to colors (and ultimately registers). A priority-based coloring algorithm [6] uses heuristics and cost analyses to determine the ordering of live ranges and guarantees that the most important live ranges are assigned colors first. In our compiler, variables (including method parameters, method local variables, class variables and compiler generated temporaries) are prioritized by their static reference counts,

having references inside loops, no matter how deeply nested, counting as 10. After the generation of the Java IR, the compiler runs data-flow analyses and performs copy propagation and common sub-expression elimination. At this point loop structures are also identified and static reference counts are calculated. The first step of our register allocator is to build a priority list of variables using this information. In case of matching static reference counts, the priority of a variable is dictated by the order in which it was declared in the code. As we want to keep the number of virtual registers as small as possible, we assign the same virtual register number to variables with non-conflicting live ranges. This is accomplished by building the interference graph which gives us information on conflicting live ranges. Using the information provided by the interference graph, the register (color) assignment algorithm picks variables from the priority list and assigns virtual register numbers (colors) to them, reusing lowest virtual register numbers or creating a new virtual register number in the case of conflicts.

In our register allocation algorithm, when assigning virtual register numbers we associate each virtual register number with the Java type of the variable it is allocated to, and we do not allow, for example, a virtual register holding an integer to later be re-used to hold a floating-point value. This restriction, although it has the counter effect of increasing the number of virtual registers, also guarantees that the mapping of a virtual register to a physical register is fixed in the run-time compiler. Otherwise, the frequent re-mapping of virtual registers to physical registers to comply with variable types and machine register assignment restrictions will conflict with the virtual register priorities, leading to an increase in spills and lower performance.

3 Annotation-Aware JIT (AJIT) Compilation System

The rightmost portion of Figure 2 depicts our annotation aware JIT (AJIT) system. We modified the public domain JIT compiler system *Kaffe* [24] (version 0.9.2) to implement our annotation scheme. The changes concentrated on a few number of files and consisted on the design of a new register allocator, modifications to the generation of *Kaffe*'s internal intermediate representation, and changes to its SPARC code generator. Both the original and new functionality coexist in the system, allowing the processing of annotated methods and non-annotated methods within the same class file.

As VRA annotations are derived from translating bytecodes into a RISC-like three address code, one wonders whether they are general, flexible and helpful enough to produce optimized code for different target architectures. We experimented with the Intel architecture

in [15], and now with the SPARC architecture in this paper — two distinct architectures (CISC and RISC respectively). Our annotation scheme has proven to suffice the needs for generating code for these two platforms. As we experiment with other architectures our annotation types and formats will be refined accordingly.

In our AJIT system, when a class method is first called, the bytecode stream is read into a table buffer. If there is an annotation code attribute, the annotations are also read into an annotations table. Then the JIT compiler invokes the corresponding translation routine. The process of producing native code from annotated Java bytecodes is done in a single pass over the bytecode stream. As each bytecode and its annotations bytes are read, the corresponding *Kaffe* IR operation(s) is (are) generated. The generated *Kaffe* IR operation (or sequence of operations) depends on the information provided by the annotations. This information may suggest that the bytecode translation be entirely skipped, or that some sub-operations be eliminated or simplified. Figure 5 shows a code example of how an iaload bytecode operation is translated using annotation information.

The translated *Kaffe* IR operation operands are specified by virtual register numbers, extracted from the annotations bytes. Once the entire bytecode stream has been processed, SPARC native code is produced from the *Kaffe* IR. At this point, as each *Kaffe* IR operation is translated into native code, the register allocator is invoked to replace virtual register numbers with machine registers.

The run-time register allocator is a fast and effective algorithm that essentially maps each virtual register to a machine register, prioritizing the assignment of lower virtual register numbers. This guarantees that high priority values (program variables represented by lower virtual register numbers) have preference in the register assignment. When the number of physical registers is exhausted, virtual registers are mapped to temporaries on the stack. In the case of the SPARC architecture, the register allocator reserves four registers of each type (four of the global integer registers g4-g7 and four of the floating point registers f28-f31) for evaluating expressions that involve variables that are not mapped into machine registers. It uses local registers l0-l7 , global registers g1-g3, any unused input register i0-i5 and floating point registers f0-f27 during allocation. Registers o0-o7 are not available for the allocator and are reserved for passing parameters to method calls. Our register allocation algorithm uses a mapping table as an auxiliary data structure. The mapping table stores information on a virtual register number, a pointer to the corresponding physical register table entry, and the stack offset value it should use

in case of spilling. There are some details on the initialization of the mapping table to correctly handle the SPARC register windows convention; these details are taken care of in the method's prologue and on the translation of bytecodes for accessing method local variables. Method local variables that are parameters are passed in special integer registers (i0-i5), forcing the mapping of virtual registers associated with these parameters.

In our experiments we observed that machine calling conventions can complicate the simple mapping-based register allocation, as it forces virtual register assignments to specific machine registers. This may break virtual register priorities, and the register allocator fixes it by spilling lower priority physical registers in case a higher priority virtual register needs a physical register and none are available. We are currently studying the effect of different calling conventions in our mapping-based dynamic register allocator.

Our current register allocation scheme does not try to minimize the cost of subroutine calls. At method call boundaries, move operations are generated to guarantee values are in the correct registers required by the calling convention and spilling of all active registers is done. Our annotation scheme could be used to carry information on which values produced in the program are later passed to methods as parameters and also which registers should be saved across procedure calls. Having the first kind of information would guide the register allocator in the virtual to physical register mapping and would avoid some copies. The second kind of information would decrease the overhead of subroutine calls by spilling only the registers that are later referenced in the program. We are currently investigating how our virtual register allocator in our annotation-generating front-end can be extended to lower the cost of method calls.

To prove that our AJIT system is an acceptable engineering solution we need to quantify the overhead of processing the annotated bytecode stream and the overhead of our mapping-based register allocation in the process of generating optimized native code. If we generate better native code more efficiently than traditional optimizing JIT compilers, we have shown that our framework is a good solution for improving the overall execution speed of Java programs. Annotation overhead results from many factors: (1) the larger class file size (which increases download time), (2) the interpretation of the information conveyed in the annotation bytes (see the extra processing required to build the *Kaffe* JIT IR in Figure 5), (3) additional work done by the run-time register allocator, and (4) the demand for extra resources (memory for storing annotations). Network applications are sensitive to the download time overhead, but other types of applications that do not depend on annotated class files being downloaded are

```
define insn(IALOAD)
{
/*
 *    ..., array ref, index -> ..., value
 */

    a = meth->annotations table->entry[i];
    i++;

    if (a.header == SRC SRC EXTRA EXTRA DEST){
        index = *(a.VRAData); objref = *(a.VRAData+1);
        tmp1 = *(a.VRAData+2); tmp2 = *(a.VRAData+3); dest = *(a.VRAData+4);

        annotated lshl int const(vrslots[tmp1].slots, vrslots[index].slots, SHIFT jint);
        if (object array offset !=0)
            annotated add int const(vrslots[tmp1].slots, vrslots[tmp1].slots, object array offset);
        annotated add ref(vrslots[tmp2].slots, vrslots[objref].slots, vrslots[tmp1].slots);
        annotated load int(vrslots[dest].slots, vrslots[tmp2].slots);

    }else if (a.header == CONST SRC EXTRA EXTRA DEST){
        cindex = *(a.VRAConst); objref = *(a.VRAData);
        tmp1 = *(a.VRAData+1);tmp2 = *(a.VRAData+2); dest = *(a.VRAData+3);

        annotated move int const(vrslots[tmp1].slots, (cindex<<SHIFT jint), NULL);
        if (object array offset !=0)
            annotated add int const(vrslots[tmp1].slots, vrslots[tmp1].slots, object array offset);
        annotated add ref(vrslots[tmp2].slots, vrslots[objref].slots, vrslots[tmp1].slots);
        annotated load int(vrslots[dest].slots, vrslots[tmp2].slots);

    }else if (a.header == SRC SRC EXTRA DEST){
        objref = *(a.VRAData); tmp1 = *(a.VRAData+1); tmp2 = *(a.VRAData+2); dest = *(a.VRAData+3);

        annotated add ref(vrslots[tmp2].slots, vrslots[objref].slots, vrslots[tmp1].slots);
        annotated load int(vrslots[dest].slots, vrslots[tmp2].slots);

    }else if (a.header == SRC DEST){
        tmp1 = *(a.VRAData); dest = *(a.VRAData+1);
        annotated load int(vrslots[dest].slots, vrslots[tmp1].slots);
    }else if (a.header == SRC){
        // no action
    } else error=1;
}
```

Figure 5: AJIT translation process for an `iload` bytecode operation

not affected. In our AJIT system, the *Kaffe* run-time IR is simple to build and manipulate. Other optimizing JIT systems will need a more complex IR to enable more advanced compiler transformations. We believe that the overhead of processing the annotations, storing them and building a simple run-time IR will ultimately be less than the overhead of building, storing and manipulating a complex IR in those systems. Finally, our run-time register allocation algorithm is an algorithm that obeys a defined mapping rule and only manipulates mapping tables. As a result, our register allocator is simple and fast. No time is spent on conflict graph construction, coloring nor dataflow analysis — tasks routinely performed by traditional register allocators.

4 Related Work

Various approaches are being proposed to overcome the inefficiency of translating the Java Bytecodes to native code, and thus increase the execution speed of Java programs. When compilation time is not a constraint, the most common approach is to translate the bytecodes into some higher-level intermediate form [7, 14] or language [21], and then back to native code (perhaps using an existing compiler, as in [21]). When speed of compilation is an issue, optimizing JIT compilers [1, 2, 16, 18, 24] try to improve the quality of the native code generated on the fly by adapting traditional optimization techniques to run-time code generation. Optimizations can also be applied during load-time, i.e. after bytecode generation yet before run-time translation to native code; [8] is an example of such a bytecode optimizer. Our annotation scheme is a hybrid approach in

that most work is done at compile-time to retain important high-level program and optimization information, while at run-time lightweight code-improving transformations accomplish the task of generating high-quality native code.

Research in the area of developing fast run-time algorithms for traditional compiler optimizations is very active. In the following paragraphs we overview commercial and academic systems, some of which make use of annotation schemes to aid code optimization. We also discuss how they implement run-time code optimizations such as common sub-expression elimination, register allocation and elimination of array bounds checking, and how these implementations compare to the run-time algorithms our annotation scheme requires. In all optimizing JIT compilers there is an attempt to develop compiler optimizations with linear time algorithms with respect to some parameter (e.g., the number of bytecode instructions, or the number of local or stack variables). Our annotation-based approach has also been designed with this in mind; our VRA annotation scheme allows run-time register allocation in linear time.

Several researchers exploit the idea of code annotations and relate to our approach. Though not designed to specifically overcome the Java Bytecode language inefficiency, these approaches could potentially be applied to this problem. In the context of dynamic code generation, code annotations in the form of programmer hints [12] or high-level language constructs extensions [20] serve as guide to where (and on what) dynamic compilation should take place. These code annotations help to build optimizing just-in-time compilers by extending to run-time the applicability of traditional compiler optimizations. Using these schemes researchers have

built different algorithms for copy propagation, dead code elimination, register allocation and even advanced cross-module optimizations. Different strategies are applied to balance the tradeoff between dynamic compilation speed and the quality of the generated code.

Most directly related to our VRA annotation scheme is the work of Wall [23] on cross-module link-time register allocation. In his approach, link-time register allocation is treated as a form of relocation. The compiler generates code that can be directly linked and executed, but it annotates some of the instructions with register actions that describe what needs to be done to the instruction if the variables it manipulates are assigned to a register at link time. Compared to our mapping-based register allocation, Wall's approach has the overhead of building the call graph and carrying out local data flow analysis at link-time, and it depends on good usage estimates (profiling information). However, it performs global register allocation while our current implementation only works intraprocedurally.

The Intel Java JIT compiler described in [2] implements a limited form of common sub-expression elimination (CSE). Our VRA annotation scheme allows a traditional CSE algorithm to be implemented at compile-time and has the further advantage in revealing common sub-expressions implicit in bytecode operations. In the Intel JIT compiler, register allocation is accomplished via a priority-based algorithm. Our mapping-based register allocation is also a priority-based scheme, but faster to implement at run-time as it dispenses with any form of code analysis. In addition, our VRA scheme can be expanded along the ideas presented in [23] to allocate global variables, while the Intel JIT compiler would need interprocedural data-flow analysis to accomplish the same, implying in an expensive run-time algorithm. A very simple array bounds check elimination algorithm was implemented in the Intel JIT compiler, handling only constant indexes. As described in [15], our run-time check annotations allow powerful subscript analysis to be performed at compile-time and easily convey this analysis information to the run-time system.

Another efficient JIT compilation system is CACAO [1]. CACAO implements copy propagation and register allocation by performing stack analysis at run-time and relying on the efficient coloring of local variables done by the Java front-end (by assigning the same local variable number to variables which are not active at the same time). CACAO also relies on the fact that stack slot variables have their lifetimes implicitly encoded. Compared to our scheme, the stack analysis information that has to be computed by their algorithm is provided for free by our VRA annotations. On the other hand, their run-time register allocator takes into account the cost of subroutine calls, which is lacking in our current scheme. If allocation of global variables is considered,

interprocedural stack analysis would be necessary, and their algorithm would become more expensive. Other optimizations such as instruction scheduling, method inlining and array bounds check removal are planned for CACAO, but the run-time cost of these additions is not clear.

Kaffe [24] is a freely available JVM that runs on several platforms; it serves as the basis for our implementation work. *Kaffe*'s native code translation process builds a simple RISC-like IR as it loops through the bytecode stream. Register allocation is combined with machine code generation. The register allocator is based on a simple algorithm that maps stack and local variable slots to machine registers. When it runs out of registers, the least recently used register is spilled and freed for allocation. There is no special treatment to reduce subroutine calling costs, or to exploit machine calling conventions, as CACAO does. Upon a call, copy operations are introduced to guarantee values are in the correct register and all modified slots are spilled. No other compiler optimizations are implemented. In Section 5 we demonstrate that our AJIT system outperforms *Kaffe* in terms of the quality of the generated native code.

An important optimizing run-time compilation system is the Slim Binary project [10, 19]. This approach proposes an architecture-neutral intermediate representation for software distribution, called *slim binaries*, that can be seen as an alternative to Java Bytecodes. The dynamic compiler for slim binaries implements code optimizations as background processes. Just like the dynamic compilation systems discussed in [12, 20], this system tries to utilize run-time information (e.g., values of variables and run-time profiling information) to perform customized optimizations. Slim binaries incorporate a more complex tree-based intermediate representation, conveying control flow information but also incurring some run-time overhead to manipulate it. Much like our annotation scheme extends the Java Bytecodes with extra information that is collected during traditional compilation, the Slim Binary representation could benefit from our annotations scheme to decrease run-time optimization costs, such as carrying extra information to aid in register allocation.

5 Results

Our results revolve around four benchmarks: Neighbor, which performs a nearest-neighbor averaging across all elements of a two-dimensional array; EM3D, a code that creates a graph and then performs a 3D electromagnetic simulation [9]; Huffman, a character string compression and decompression application; and Bitonic Sort, which builds a binary tree and then performs bitonic sorting (recursively) [3]. To measure the impact

of our AJIT system, we collected results using JVMs available on the SPARC platform: Sun's JDK version 1.1.1 [17] and *Kaffe* JVM version 0.9.2 [24]. The execution time results are shown in Table 1. Note that the timings do not include translation nor compile time, and thus represent the quality of the generated code. All codes were compiled using our annotation-generating Java Bytecode compiler and then executed using Sun's interpreter, the *Kaffe* JIT compiler, and our AJIT system.

Benchmarks	SUN Interpreter (in secs)	kaffe JIT (in secs)	AJIT (in secs)
Neighbor 256X256 array Iterations = 1500	553.03	162.73	115.31
EM3D 1250 tree nodes Iterations = 200	359.84	149.86	74.51
Bitonic Sort 1024 tree nodes Iterations = 512	167.05	141.23	120.96
Huffman 30000 array nodes Iterations = 288	4690.00	1856.00	1487.00

Table 1: Benchmarks execution times (in seconds)

Benchmarks	SpeedUp AJIT/SUN	SpeedUp AJIT/Kaffe
Neighbor	4.80	1.41
EM3D	4.83	2.01
Bitonic Sort	1.38	1.17
Huffman	3.15	1.25

Table 2: Benchmarks speedups

The results presented in Table 1 reflect the sole effect of our VRA annotations scheme. From the two speedup columns of Table 2 we see that our annotation based approach offers speedups varying from 1.38 to 4.83 over direct interpretation, and is 17% to 100% faster than *Kaffe*'s JIT technology. Also notice that the best speedups were achieved for codes consisting of basic loops iterating over array-based or pointer-based data (Neighbor, EM3D and Huffman). For such codes, the VRA annotations helped to identify common subexpressions and eliminate them, along with the propagation of values and elimination of move operations. These transformations correspond to optimizations that could not be expressed in the Java Bytecodes directly. The annotations also ensured that the most important

variables, such as loop index variables, were permanently assigned to machine registers throughout method execution. The smallest performance gain was observed for the code with the highest number of subroutine calls — Bitonic Sort, a recursive algorithm. This result is explained by the way our AJIT system, and *Kaffe* as well, handle subroutine calls during dynamic register allocation. In short, both JIT compilers do not take advantage of SPARC register windows. All active registers are saved across method calls, introducing significant overhead. Thus, it is important to note that this overhead is not an intrinsic limitation of the algorithms, but an artifact of the current implementations.

The encouraging observation we have obtained from these preliminary results is that despite many limitations of the first implementation, our AJIT system is capable of producing machine code that executes up to twice as fast as current JIT technology. By extending our VRA annotations scheme with extra information, such as registers to be saved across subroutine calls, and improving the implementation of the dynamic register allocator itself, we believe the impact of our VRA annotations will be even more significant.

6   Conclusions and Future Work

Most approaches for speeding up Java execution resort to dynamic compilation (and even dynamic code re-optimization [13]). In this scenario, run-time costs must be minimized and thus it is desirable that the bulk of the compilation process be done statically at compile time. Having a rich program representation conveying, for example, dependence information to allow instruction scheduling and support for dynamic register allocation, will decrease the time spent on run-time code generation by cutting down the time spent on program analysis and transformation. In this paper we discussed how the Java Bytecode language is a poor choice for a high-performance program representation, since it demands a more time consuming code generation process (at run-time!) in order to produce high-quality native code. We presented an approach based on code annotations that helps overcome this problem, and discussed the implementation details of our resulting annotation-aware JIT system.

Our first prototype implements the VRA annotation scheme that conveys information for dynamic register allocation. It also enables some basic code scheduling by identifying and eliminating redundant computation and allowing propagation of values. Preliminary results show that we outperform JIT technology, producing code that runs up to twice as fast. We plan to extend our VRA annotation scheme by incorporating information that helps minimize the cost of subroutine calls (e.g., values to be saved across procedure calls and val-

ues passed as subroutine parameters) and allows cross-module register allocation. We started with the implementation of the VRA annotations scheme because register allocation is the most important compiler optimization on today's architectures. We also initially selected scientific benchmarks to test our approach given their higher sensitivity to such optimization. Figure 2 mentions a number of annotation possibilities that we plan to explore in the future. These annotations support more sophisticated compiler optimizations, such as instruction scheduling and lifetime analysis for reducing garbage collection. To help evaluate these annotations we will study non-numeric Java benchmarks as well.

References

[1] R. Graft A. Krall. Efficient JavaVM Just-in-Time Compilation. *In Proceedings of International Conference on Parallel Architectures and Compilation Techniques, PACT'98*, 1998.

[2] A. Adl-Tabatabai, M. Cierniak, G. Lueh, V. M. Parikh, and J. M. Stichnoth. Fast, Effective Code Generation in a Just-In-Time Java Compiler. *Proceedings of ACM Programming Languages Design and Implementation*, pages 280–290, 1998.

[3] G. Bilardi and A. Nicolau. Adaptive Bitonic Sorting: An Optimal Parallel Algorithm for Shared Memory Machines. Technical Report TR86-769, Cornell University, 1986.

[4] G. J. Chaitin. Register Allocation and Spilling via Graph Coloring. *SIGPLAN Notices*, 17(6):201–107, June 1982.

[5] G. J. Chaitin, M. A. Auslander, A. K. Chandra, J. Cocke, M. E. Hopkins, and P. W. Markstein. Register Allocation via Coloring. *Computer Languages*, 6:47–57, January 1981.

[6] F. C. Chow and J. L. Hennessy. A Priority-based Coloring Approach to Register Allocation. *ACM TOPLAS*, 12(4):501–536, October 1990.

[7] M. Cierniak and W. Li. Optimizing Java Bytecodes. *Concurrency: Practice and Experience*, 9(11), November 1997.

[8] L. R. Clausen. A Java Bytecode Optimizer Using Side-effect Analysis. *Concurrency: Practice and Experience*, 9(11), November 1997.

[9] D. Culler, A. Dusseau, S. Goldstein, A. Krishnamurthy, S. Lumetta, T. von Eicken, and K. Yelick. Parallel Programming in Split-C. In *Proceedings of Supercomputing 1993*, pages 262–273, November 1993.

[10] M. Franz and T. Kistler. Slim Binaries. *Communications of the ACM*, 40(12):87–94, December 1997.

[11] J. Gosling, Bill joy, and G. Steele. The Java Language Specification. Addison-Wesley, 1996.

[12] B. Grant, M. Mock, M. Philipose, C. Chambers, and S. J. Eggers. Annotation-Directed Run-Time Specialization in C. *In Proc. of PEPM*, June 1997.

[13] David Griswold. The Java HotSpot Virtual Machine Architecture, March 1998. See whitepaper at http://www.javasoft.com/products/hotspot/.

[14] C. Hsieh, J. Gyllenhaal, and W. Hwu. Java Bytecode to Native Code Translation: The Caffeine Prototype and Preliminary Results. *Proceedings of the 29th Annual Workshop on Microprogramming*, December 1996.

[15] J. Hummel, A. Azevedo, D. Kolson, and A. Nicolau. Annotating the Java Bytecodes in Support of Optimization. *Concurrency: Practice and Experience*, 9(11):1003–1016, November 1997.

[16] Microsoft Inc. The Microsoft Virtual Machine for Java.
See http://www.microsoft.com/java/sdk/.

[17] SUN Inc. Sun interpreter.
See http://www.javasoft.com.

[18] Symantec Inc. Just in Time Compiler for Windows 95/NT. See http://www.symantec.com.

[19] T. Kistler and M. Franz. Dynamic Runtime Optimization. *In Proceedings of the Joint Modular Languages Conference, JMLC'97*, pages 53–66, March 1997.

[20] M. Poletto, D. R. Engler, and M. F. Kaashoek. tcc: A System for Fast, Flexible and High-Level Dynamic Code Generation. *Proceedings of ACM Programming Languages Design and Implementation*, 1997.

[21] T. Proebsting, J. Hartman, G. Townsend, P. Bridges, T. Newsham, and S. Watterson. Toba: A Java-to-C translator.
See http://www.cs.arizona.edu/sumatra/toba.

[22] Effective Edge Technologies. guavac.
See summit.stanford.edu:/pub/guavac/.

[23] D. W. Wall. Global Register Allocation at Link-Time. *In Proc. ACM SIGPLAN'86 Symp. on Compiler Construction*, pages 264–275, June 1986.

[24] Tim Wilkinson. Kaffe: A Free JIT virtual machine to run Java code.
See http://www.transvirtual.com.

TC2100 Bldg./Room RANDOLPH

Organization

U. S. DEPARTMENT OF COMMERCE

COMMISSIONER FOR PATENTS

P.O. BOX 1450

ALEXANDRIA, VA 22313-1450

IF UNDELIVERABLE RETURN IN TEN DAYS

OFFICIAL BUSINESS

AN EQUAL OPPORTUNITY EMPLOYER

UNDELIVERABLE AS
ADDRESSED
FORWARDING ORDER
EXPIRED

RECEIVED

JUN 2 0 2005

USPTO MAIL CENTER